

- 1 -

16670IDWO02W: FRISKNEY, BRAGG, PARRY, ASHWOOD-SMITH,
ALLAN, BRUECKHEIMER, GIBBON

**DIFFERENTIAL FORWARDING IN ADDRESS-BASED CARRIER
NETWORKS**

5

FIELD OF THE INVENTION

The present invention relates to methods of, software for and apparatus for differential forwarding in address-based carrier networks. In particular, but not exclusively, the present invention relates to methods of, software for and apparatus for differential forwarding and establishing connections in intrinsically connectionless carrier networks which address-based forwarding, such as Ethernet or IP networks.

BACKGROUND TO THE INVENTION

15 For many years now, telecommunications carriers have been deploying packet-switched networks in place of or overlaid upon circuit-switched networks for reasons of efficiency and economy. Packet-switched networks such as Internet Protocol (IP) or Ethernet networks are intrinsically connectionless in nature and as a result suffer from Quality of Service (QoS) problems. Customers value services which are guaranteed in terms of bandwidth and QoS.

20 Carriers may use Multi-Protocol Label Switching (MPLS) over a layer 2 network to create connection-oriented label switched paths (or tunnels) across the intrinsically connectionless network, and thereby to provide guaranteed QoS and bandwidth services to customers. However, MPLS is a

- 2 -

relatively unstable and complex standard and carriers ideally desire an alternative.

It is desired to use Ethernet switches in carriers' networks. Use of Ethernet switches in carriers' networks would have the advantages of interoperability (mappings between Ethernet and other frame/packet/cell data structures such as IP, Frame Relay and ATM are well known) and economy (Ethernet switches are relatively inexpensive compared to IP routers, for example). It would also provide a distinct advantage of being the principal technology used by enterprises that require a wide area network service from a carrier and therefore able to work in a native mode.

However, the behaviour of conventional switched Ethernet networks is incompatible with carriers' requirements for providing guaranteed services to customers. Carriers need networks to be meshed for load balancing and resiliency – ie there must be multiple paths across it – and the ability to perform traffic engineering – ie the ability of the network operator to control the provision of explicitly routed variable bandwidth connections (or tunnels) through which traffic may be directed. This provides operators significant flexibility in that the physical network build is not obliged to correspond to the offered load and therefore is tolerant of changing usage patterns without requiring on going physical modifications.

In contrast, conventional Ethernet networks must be simply-connected - ie there must be one and only one logical path choice between each and every node of the network. As a consequence, conventional Ethernet networks do not have support for network-wide load balancing, suffer from resiliency problems and cannot support traffic engineering. Further the impact of a single failure with respect to the overall load carried can be significant.

Spanning tree protocols are known which enable a physically meshed Ethernet network to be logically transformed into a simply-connected network by detecting physical loops and logically disabling connections to break up the loops. Spanning tree protocols are also known which are able to detect

- 3 -

failure of a physical connection (thereby partitioning the fully-connected network) and automatically restore one or more previously-disabled physical connections so as to re-connect the network. This provides a degree of resiliency. However, carriers need to plan their network traffic routes to achieve much higher resiliency, flexibility and efficiency than spanning tree can achieve. This level of routing capability is best achieved by segregating the traffic into connections whose routes are determined as part of this planning process.

Virtual Bridged LANs (or VLANs) are described in the Institute of Electrical and Electronics Engineers (IEEE) standard 802.1Q, 2003 Edition. Figure 1 shows a conventional VLAN 10 split up into a plurality of component LANs 12 and connected via VLAN-aware Media Access Control (MAC) bridges 14. Component LANs 12 are typically provided for different communities of interest, such as users sharing a common server or having common network protocol requirements. Unique identifiers (VLAN tags or VLAN IDs) are used to identify each component LAN. Broadcast traffic is broadcast only within component LANs. This helps to overcome the scalability issues of Ethernet by partitioning the whole network 10 resources into smaller broadcast domains. VLAN tags are used to distinguish between traffic for different component LANs when forwarding traffic on shared links between MAC bridges. However the size of the standard VLAN tag is limited to 12 bits, which in turn limits the scale of the network and the number of partitions of component LANs to 4094, where two VLAN tags are reserved with special meaning not for general assignment.

The Internet Engineering Task Force (IETF) has published an Internet Draft referred to as draft-kawakami-mpls-lsp-vlan-00.txt. This document describes the use of VLAN tags for label switching across Ethernet networks in a manner similar to use of MPLS labels for label switching over MPLS networks - VLAN tags are used as labels to mark traffic at an ingress point of a label switched path (LSP) as belonging to a Layer 2 tunnel, and VLAN-aware Ethernet switches in the network act as a VLAN label switched routers.

- 4 -

Connections are formed using one or more LSPs. Intermediate nodes along the connection may optionally swap the inbound label to a different outbound label. In this manner the VLAN tag has meaning specific to any given local node, and the ability to reuse VLAN tags solves some of the scalability issues of 802.1Q.

However, one problem with the method proposed in draft-kawakami-mpls-lsp-vlan-00.txt is that only a maximum of 4094 unique VLAN tags are definable in 802.1Q compliant equipment. This still limits the flexibility and increases the complexity of provisioning connections across the network. Another problem is that connections may not easily be re-routed once provisioned without in general creating transitory loops.

Another problem is that since the Frame Check Sequence (FCS) in Ethernet frames is computed over both the payload and header portions of the frame, every time a VLAN tag (ie a label) is swapped at the ingress or egress point of a LSP, the FCS needs to be recomputed since the VLAN tag will have changed. This requires performing a computation function over the entire Ethernet frame. Moreover, during the interval from when the original FCS is removed and the new FCS added, the frame is vulnerable to corruption without the protection of any FCS.

Yet another problem with the ‘label-swapping’ approach proposed in draft-kawakami-mpls-lsp-vlan-00.txt is that it requires a “chain of correctness”, in that forwarding relies on each local label-forwarded link on the LSP being correct. This should be contrasted with conventional Ethernet which uses globally unique address information to perform forwarding. As the LSP labels are not globally unique per conventional Ethernet, it is possible for a forwarding fault in performing label translation to be concealed if a value is incorrectly mapped to another value that is in use. More importantly, from a practical perspective, ‘label-swapping’ behaviour represents a significant change from conventional Ethernet switch functionality, and current telecommunications standards.

SUMMARY OF THE INVENTION

The present invention relates to enabling the establishment of connections in frame-based networks such as Ethernet networks. The capability of establishing connections in an Ethernet network provides the ability to partition the network resources in a specific way that could be for traffic engineering purposes, to pin paths over physically diverse routes for improved resilience, and monitor those resources for performance purposes, audit integrity, locate faults and other proactive verification purposes. The term traffic engineering is used broadly in the present document to refer to functions for establishing and maintaining the quality of service of the customers' connections while permitting the owner to operate their network efficiently. Examples of this are ensuring that no link is over-loaded, load-balancing the connections in a preferred manner across the network, re-planning the load on the network by re-routing some existing connections, establishing protection mechanisms, performing traffic restoration actions, relative priority of different traffic types, admission control, policing, scheduling and so on.

According to an embodiment of the present invention, connections are established in the carrier network by configuring, in one or more network nodes, mappings for forwarding data frames such as Ethernet frames. The mappings are from a combination of a) a destination (or source) address corresponding to a destination (or source) node of a connection, such as a MAC address, and b) an identifier, which for Ethernet may be a reuse of a VLAN tag which then is no longer necessarily unique to a subnet, but only necessarily unique when combined with the destination (or source) address. The mappings are to selected output ports of the one or more nodes. By using the combination of address AND identifier, the mappings enable data frames belonging to different connections to be forwarded differentially (ie forwarded on different output ports) despite the different connections

- 6 -

potentially having the same destination (or source) node. This enables flexibility in routing connections – eg the ability to perform traffic engineering. The reader should note that the term address is used in this document to denote any means of identifying a network node or an ingress or egress interface of a network node, or any sub unit of a network node, for example a port card or an encapsulation function of a network node.

According to a first aspect of the present invention, there is provided a method of establishing a connection in a network, the method comprising the step of:

10 configuring, in a node of the network, a first mapping for use in forwarding data frames, the first mapping being from a combination of:

a first network address uniquely identifying, within an addressing scheme of the network, a first node of the network, and

a first identifier,

15 the first mapping being to a selected output port of the node,

the configuring thereby establishing at least part of a first connection for forwarding data frames, the connection being through the node,

the first identifier being a qualifier of the first network address, the combination thereby enabling differential forwarding, at the node, of data frames addressed to or from the first node.

According to a second aspect of the present invention, there is provided a connection controller for establishing a connection in a network, the connection controller comprising:

25 a signal generator arranged in use to generate a first signal for configuring, in a node of the network, a first mapping for use in forwarding data frames, the first mapping being from a combination of:

- 7 -

a first network address uniquely identifying, within an addressing scheme of the network, a first node of the network, and

a first identifier,

the first mapping being to a selected output port of the node,

5 the configuring thereby establishing at least part of a first connection for forwarding data frames, the connection being through the node,

the first identifier being a qualifier of the first network address, the combination thereby enabling differential forwarding, at the node, of data frames addressed to or from the first node.

10 According to a third aspect of the present invention, there is provided a network comprising a node configured with a first mapping for use in forwarding data frames, the first mapping being from a combination of:

a first network address uniquely identifying, within an addressing scheme of the network, a first node of the network, and

15 a first identifier,

the first mapping being to a selected output port of the node,

the configuring thereby establishing at least part of a first connection for forwarding data frames, the connection being through the node,

the first identifier being a qualifier of the first network address, the combination thereby enabling differential forwarding, at the node, of data frames addressed to or from the first node.

20 A communications network comprising one or more nodes arranged to perform the method of the first aspect of the present invention set out above is also provided.

- 8 -

A computer program arranged to perform the method of the first aspect of the present invention set out above is also provided.

Advantageously, the present invention enables connections to be established in a frame-based network in a highly flexible manner enabling network-wide traffic engineering. Furthermore, the specific problems inherent in the method proposed in draft-kawakami-mpls-lsp-vlan-00.txt (processing overhead and vulnerability of frames to corruption) are overcome since no label swapping is performed.

According to a fourth aspect of the present invention, there is provided a method of establishing connections in a frame-based network, the method comprising the step of configuring, in one or more nodes of the network, first mappings for use in forwarding data frames, the first mappings being from a combination of a first destination address corresponding to a first destination node of the network, and a first identifier, the first mappings being to a selected output port of, or to respective selected output ports of each of, the one or more nodes, thereby establishing at least part of a first connection through the one or more nodes to the first destination node.

In one embodiment, the method of the present invention includes configuring, in at least one of the nodes, a second mapping for use in forwarding data frames, the second mapping being from a combination of: a second destination address corresponding to a second destination node of the network, and a second identifier, the second mapping being to a selected output port of the at least one node, thereby establishing at least part of a second connection through the at least one node to the second destination node, the selected output ports of the at least one node being different for the first and second mappings, thereby enabling, at the at least one node, differential forwarding of data frames associated with the first and second connections.

Thus, advantageously, two connections may be established which converge in route at an intermediate node and then diverge again, for example.

- 9 -

In one embodiment, the first and second destination addresses and the first and second destination nodes are the same. Thus, for example, two connections may be established which converge at an intermediate node and then diverge, despite having the same destination node. This enables greater flexibility in setting up connections.

In one embodiment, the first and second identifiers are the same. Thus, for example, two connections may be established which converge at an intermediate node or nodes and then diverge, despite using the same identifier. Thus, limitations on the number of values identifiers can take do not significantly reduce flexibility in traffic engineering.

Preferably, the network is an Ethernet network and the one or more nodes are Ethernet switches. Preferably, the identifier is a VLAN tag. Advantageously, this enables traffic engineered carrier networks to be deployed using conventional and relatively inexpensive VLAN-aware Ethernet switches, albeit configured in an entirely novel and inventive manner.

In one embodiment, the configuration is performed by a control plane of the network. Thus, carriers have direct control over the establishment of traffic engineering connections in the network. Preferably, the control plane is ASON/ASTN. The control plane may be centralised or distributed.

A frame-based communications network comprising one or more nodes arranged to perform the method of the first aspect of the present invention set out above is also provided.

Software arranged to perform the method of the first aspect of the present invention set out above is also provided.

According to a fifth aspect of the present invention, there is provided a connection controller for establishing connections in a frame-based network, the connection controller comprising: a signal generator capable of generating a first signal for configuring, in a transport node of the network, a first

- 10 -

mapping for use in forwarding data frames, the first mapping being from a combination of: a first destination address corresponding to a first destination node of the network, and a first identifier, the first mapping being to a selected output port of the transport node, the first signal thereby establishing at least part of a first connection through the transport node to the first destination node.

According to a sixth aspect of the present invention, there is provided a method of establishing a connection in a frame-based network, the method comprising the steps of: configuring forwarding information in a plurality of nodes of the network the forwarding information enabling the nodes to forward data frames in dependence on a combination of a destination address and an identifier of the data frames.

According to a seventh aspect of the present invention, there is provided a method of data traffic engineering in a frame-based network, the method comprising the following steps: establishing a first and second connections in the network passing through a common switching node of the network, configuring the switching node to forward data frames differently in dependence on differences in either a destination address or an identifier of the data frames, thereby enabling data traffic engineering.

According to a eighth aspect of the present invention, there is provided a method of establishing connections in a frame-based network, the method comprising the step of: configuring, in each of a first plurality of nodes of the network, a first forwarding mapping from a first combination of a destination address and an identifier to a selected output port of a respective node of the first plurality of nodes.

According to a ninth aspect of the present invention, there is provided a connection controller for establishing connections in a frame-based network, the connection controller being arranged to configure a first forwarding mapping in a transport node, the first mapping being from a first combination

- 11 -

of a destination address and an identifier to a first output port of the transport node.

According to a tenth aspect of the present invention, there is provided a method of forwarding data frames in a frame-based network, the method comprising the steps of: establishing a first connection in the network, the first connection being associated with a first combination of a destination address and an identifier, and forwarding data frames in the network in dependence on a combination of a destination address and an identifier of the data frames.

Further aspects of the present invention are set out in the appended claims. Further advantages of the present invention will be apparent from the following description.

In order to show how the invention may be carried into effect, embodiments of the invention will now be described by way of example only and with reference to the accompanying figures in which:

BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 shows a conventional Virtual Bridged LAN;

Figure 2 shows an arrangement of Ethernet switches forming a carrier network according to the present invention;

Figure 3 shows a control plane/transport plane architecture for controlling the Ethernet carrier network of Figure 1 according to the present invention;

Figure 4 shows the carrier Ethernet network of Figure 1 arranged to provide connectivity between customer sites according to the present invention;

- 12 -

Figure 5 shows how nodes of the control plane interact with Ethernet switches of the transport plane to establish a connection across carrier network according to the present invention;

5 Figure 6 is a flow diagram showing the preferred use of VLAN tag and destination address to differentiate forwarding of data traffic in different connections across the carrier network, according to the present invention;

Figure 7 shows an example of differential forwarding for two traffic flows having the same source and destination provider edge nodes but different VLAN tags according to the present invention;

10 Figure 8 shows an example of differential forwarding for two traffic flows having the same source provider edge nodes and VLAN tags but different destination provider edge nodes according to the present invention;

15 Figure 9 shows an example of converged routing for two traffic flows having the same destination provider edge node and VLAN tags but different source provider edge node according to the present invention;

Figure 10 shows a sparse mode of broadcast operation for customer VPNs provisioned across a carrier network, according to the present invention;

Figure 11 shows a dense mode of broadcast operation for customer VPNs provisioned across a carrier network, according to the present invention; and

20 Figures 12 to 14 show arrangements for providing a Virtual Private LAN Service (VPLS) according to the present invention.

DETAILED DESCRIPTION OF INVENTION

25 Embodiments of the present invention are described below by way of example only. These examples represent the best ways of putting the

- 13 -

invention into practice that are currently known to the Applicant although they are not the only ways in which this could be achieved.

To support guaranteed QoS to customers, what is required is:

- 1) an at least partially meshed carrier network;
- 5 2) the ability to establish explicitly routed connections across the carrier network between any two edge nodes (traffic engineering); and
- 3) the ability to enforce any bandwidth restrictions and/or forwarding rules applied to the connections.

The present invention is primarily concerned with enabling requirements 1)
10 and 2) above in frame-based networks such as Ethernet networks.
Requirement 3) may be achieved for example using conventional mechanisms such as admission control in either or both of the control plane and at the ingress nodes of connections (trusted-edge policing). Alternatives to achieving requirement 3) are set out later in more detail.

15 Figure 2 shows an arrangement of Ethernet switches and communications links forming a carrier network according to the present invention. Carrier network cloud 20 comprises Ethernet switches 22a, 22b, 24a, 24b, 26 and 28. Ethernet switches 22a, 22b and 26 are located at the edges of carrier network 20, whereas Ethernet switches 24a, 24b, and 28 are located in the core network. Communications links (shown as straight lines in Figure 2) are provided between Ethernet switches 22a, 22b, 24a, 24b, 26 and 28. These communications links may be for example relatively long distance links over optical equipment such as SONET/SDH equipment with Ethernet interfaces using Generic Framing Procedure (GFP) (ITU-T Recommendation
20 G.7041/Y.1303).

25

Note that core network switches 24a, 24b, and 28 are fully-meshed – ie there is a direct communications link connecting each core network switch 24a, 24b, and 28 to each other. Edge network switches 22a, 22b and 26 are not

- 14 -

fully-meshed but have at least one direct communication link to communications link to a core network switch 24a, b or 28. The reader will appreciate that the particular network arrangement described is exemplary. In general, carrier networks may be implemented with virtually any number of Ethernet switches which, according to the present invention, may be connected in a fully-meshed or partially-meshed manner.

Figure 4 shows how a carrier Ethernet network may provide connectivity between customer sites according to the present invention. Three customers having respective pairs of geographically distant Ethernet switches (40a and 40b, 42a and 42b, and 44a and 44b) are shown connected to carrier network 20 via edge Ethernet switches 22a and 22b respectively. The communications links between edge switches 22a and 22b and customer switches 40a, 40b, 42a, 42b, 44a, and 44b may be dedicated links such as T1, E1 leased lines or access links such as digital Subscriber Lines (DSLs).

Carrier edge switches 22a, 22b (and 26 in Figure 2) may be logically separated into a single Provider Edge- (PE-) Core and one or more PE-Edge functions. The PE-Edge is the ingress/egress point at which customer traffic enters or leaves the provider network - ie carrier network 20. The PE-Core preferentially encapsulates incoming Ethernet traffic from the customer using Media Access Control (MAC) in MAC encapsulation (or if desired Pseudo-Wire over MAC encapsulation) and forwards the encapsulated traffic across the carrier network. This embodiment is preferred as a mechanism to limit the number of table entries required because only the MAC address space of the carrier network need be recognised, and not the whole of customer's MAC address space which could be changed independently. Similarly the PE-Core decapsulates (strips) outgoing Ethernet traffic and forwards the stripped traffic on to the customer via the appropriate PE-Edge. VLAN tags are used to provide customer separation at the logical PE-Core with each different customer site connected to each edge switch having a unique VLAN tag. Stacked VLAN (ie VLAN in VLAN encapsulation or Q-in-Q) may be used to protect any VLAN tags used by the customer traffic.

- 15 -

For example, customer switch 42a may send Ethernet traffic over communications link 46a to the logical PE-Edge of edge switch 22a. Logical PE-Core of edge switch 22a encapsulates each Ethernet frame in a further Ethernet frame using the MAC address of edge switch 22a as the source address and the MAC address of the appropriate egress point – in this case edge switch 22b – as the destination address. The encapsulated traffic is forwarded across a connection established over communications links 48 of carrier network 20 to edge switch 22b. Connections may be typically trunked in the sense that traffic from multiple customers will be routed through the same connection. Alternatively, those skilled in the art will appreciate that separate connections 48 could be used for each customer. At the PE-Core of edge switch 22b, the original frames are stripped of their encapsulation and sent over communications link 46b via PE-Edge of edge switch 22b to customer switch 42b.

15

The reader will appreciate that in alternative embodiments of the present invention the logical PE-Edge may also be physically separated from the logical PE-Core and may reside at customer premises whereas the PE-Core would preferentially reside at a central office or Point of Presence (PoP) of the carrier. The reader will also appreciate that other edge switches 26 (Figure 2) may also have connections to customer sites and that customers may have to be provided with connectivity between two or more geographically distant sites over carrier network 20.

25

It will now be described how carrier network 20 is arranged to establish connections through which to forward encapsulated Ethernet traffic. A connection may be defined as an entity configured in a network which provides transport of data from a source node to one or more sink nodes.

30

As described above, carrier network 20 must be at least partially-meshed – ie there must be more than one possible path between at least some, and preferably all, nodes of the network. Thus, as will be explained below,

- 16 -

Ethernet MAC address auto learning functionality should preferably be at least partially deactivated.

On start-up (or on re-start), conventional switched Ethernet networks behave like a "classic" Ethernet Local Area Networks (LANs) in that every Ethernet frame is broadcast across the entire network. Thus, every switch, receiving an Ethernet frame on one port, broadcasts the frame out on every other port. The process repeats as the frame is received by other switches thus broadcasting the frame across the entire network.

MAC address auto-learning functionality is provided to improve configuration efficiency in switched Ethernet networks. Ethernet frames have source and destination MAC addresses corresponding to their source and destination Ethernet switches. (An Ethernet switch here is signifying an end system which is therefore configured with at least one MAC address.) When an Ethernet frame sent out by a source switch is received by an intermediate or destination Ethernet switch, the receiving switch observes the port on which the frame was received and the source address of the frame. It then builds up a forwarding table for use in future frame switching. The forwarding table maps destination address to output port and is built up using the source address of a received frame and the input port on which it was received. Over time, the network builds up forwarding state enabling efficient switching of Ethernet frames, without relying on broadcast any further.

It can thus be seen that conventional switched Ethernet networks using auto-learning must be simply-connected - ie there must be one and only one path between each and every node of the network. If there were multiple paths between any two nodes, the input port on which a frame is received from a source node would not be a reliable indicator of the correct output port to forward future traffic destined for that node. Inconsistencies in forwarding tables on Ethernet switches could result in looping of frames. Moreover, if there exists any loop in a part of the network then any broadcast packet will be continuously duplicated in that loop and the duplicates forwarded all over

- 17 -

the whole network, limited only by the link capacities concerned. This inevitably results in catastrophic failure of the network.

According to the present invention, instead of using auto learning to configure forwarding tables in Ethernet switches, forwarding tables are directly configured using a novel Ethernet control plane. Figure 3 shows a control plane/transport plane architecture for controlling the Ethernet carrier network of Figure 1. The ITU-T Automatically Switched Transport Network (ASTN), sometimes known as the Automatically Switched Optical Network (ASON), may be used preferentially. The general architectural specification of the ASTN is set out in ITU-T Recommendation G.8080.

Control plane 30 comprises a number of connection controllers 32a, 32b, 34a, 34b, 36 and 38 corresponding to each of Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20 (not all connection controllers are labelled in Figure 3, for clarity). Control Plane 30 may be conceptually thought of as lying 'above' transport plane 32 which comprises the Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20. Connection controllers (CCs) 30 are logical agents each corresponding to a respective Ethernet switch (which represent cross connects in ASTN terminology) in transport plane 32. Each CC controls the switching of its respective switch using Connection Control Interface (CCI) signalling (shown as dotted lines in Figure 3). CCI signalling is used to directly configure the forwarding tables used by Ethernet switches 22a, 22b, 24a, 24b, 26 and 28 of carrier network 20. CCs may communicate between themselves using a Network to Network Interface (NNI). Typically, CCs will exchange information regarding their operational state and the state, in particular the capacity, of their communications links using NNI signalling. Other control plane functions such as heartbeat, ping and circuit monitoring may be provided using the ITU-T standard-in-preparation currently referred to as Y.17ethOAM or the methods in IEEE standard 802.1ag.

- 18 -

While CCs 32a, 32b, 34a, 34b, 36 and 38 are logically separate from Ethernet switches 22a, 22b, 24a, 24b, 26 and 28, the reader will understand that they may be implemented in the same physical nodes in a distributed control plane model. Additionally, one CC may control one or more Ethernet switches which is moving towards a more centralised control plane model.
5 Furthermore, NNI signalling may take place over the same communications links used for transporting user traffic.

Figure 5 shows how control plane 30 interacts with transport plane 32 to establish a point-to-point connection across carrier network 20. Typically, the
10 connection will be bi-directional, although this can simply be considered as the combination of two uni-directional point to point connections. A request to establish a connection specifying a required bandwidth and an explicit route across carrier network 20 is generated for example by a supervisory network management node (not shown) or distributed network management system or function.
15 The explicit route will have been determined in accordance with a conventional routing protocol taking into account the topology of the carrier network, the operational state of network resources and the bandwidth requirements of existing and possible future connections. The route to be taken by the exemplary connection shown in Figure 5 spans Ethernet switches 22a, 24a, 24b and 22b over communications links 48. Since the
20 connections share many qualities with SONET/SDH trails, management systems already developed for SONET/SDH trail management may be re-used for managing connections in the present invention – for example Nortel's Optical Network Manager. This has the advantage that carrier networks already using SONET/SDH trail management systems need not invest in new management systems when deploying the network arrangements proposed in the present invention. The route may also be established by direct NNI signalling between CCs in an ad-hoc fashion.
25

The request to establish a connection is first sent to CC 32a. On receipt of
30 the request, CC 32a checks whether the communications link between switches 22a and 24a has sufficient capacity to support the required

- 19 -

bandwidth. If so, it forwards a connection setup request message 50 to CC 34a specifying the required bandwidth and explicit route. CC 34a then checks whether the communications link between switches 24a and 24b has sufficient capacity to support the required bandwidth. The process continues
5 until the connection setup message request 50 reaches CC 32b. Along the route, CCs may optionally reserve bandwidth of their respective switches and communication links so as to avoid race conditions where competing connections are setup over the same resources.

When connection setup request message 50 reaches CC 32b, if there is
10 sufficient bandwidth along the entire path to support the required connection, then CC 32b sends a connection setup response message 52 back to CC 34b, CC 34a and finally to CC 32a. As the connection setup response message 52 traverses the CCs, each CC sends CCI signalling 54 to its respective switch to configure the forwarding tables of each switch, thereby to
15 establish the forwarding state required to setup the connection.

It will be appreciated that the mechanism for establishing connections across carrier network 20 described above is merely exemplary and other well-known mechanisms may be used. For example, all the admission control may be performed in a centralised CC controlling several if not all the Ethernet
20 switches in the extreme. In another example arrangement, the supervisory management function may be used to compute routes for connections and simultaneously perform the necessary admission control; this in turn would simplify the role performed by the CC. Yet another example is where the supervisory management function or CC consults a specialised and either
25 centralised or distributed Bandwidth Manager or Policy Decision Function to perform the admission control.

How forwarding tables of the Ethernet switches of carrier network 20 are used to support connections is a key aspect of the present invention and will now be described in detail.

- 20 -

Typically, there will be many thousands or tens of thousands of connections established across a carrier network at any time. These connections will share the physical resources of the carrier network – ie the switches and communications links. Thus, each switch will typically have a large number of 5 connections established through it at any point in time. However, each switch must be able to forward data traffic according to the explicit route requirements of the specific connection through which that traffic is being sent. A likely scenario is that the carrier network will need to establish multiple connections from the same source nodes, multiple connections to the 10 same destination nodes and multiple connections both from the same source nodes and to the same destination nodes. However, for traffic engineering purposes, the latter connections may need to be established through physically distinct routes across the network. Furthermore, these routes may need to converge and diverge again within the carrier network. To support 15 such route flexibility in connections, what is required is that each switch be able to differentiate between data traffic travelling in different connections and forward accordingly.

However, conventional switched Ethernet is incapable of this. As described above, conventional Ethernet switches forward traffic based solely on a 20 forwarding table (established through auto learning) mapping destination address to output port. As a result, a conventional Ethernet switch will not be able to differentiate between data traffic having the same destination address, although it may be associated with multiple different connections.

According to the present invention, VLAN tags are used to enable 25 differentiation of connections established across a carrier network and thereby to enable differential forwarding. Preferentially the VLAN tag defined in IEEE 802.1Q is applied in a novel manner such that the Ethernet switches of carrier network 20 are 802.1Q VLAN-aware but arranged to use a combination of destination address and VLAN tag to forward data traffic. This 30 is preferentially achieved by reusing the existing capabilities in each Ethernet switch to store separate forwarding tables for each VLAN tag configured, the

- 21 -

VLAN tag acting as a mapping (or indexing) to forwarding tables, and each forwarding table mapping destination address to output port. However in the present invention the group of forwarding tables are arranged to provide a mapping from a combination of destination address and VLAN tag to output port.

5

According to the preferred embodiment, VLAN tags have meaning only within the context of a destination address. As such, the allocation of VLAN tags is logically localised to the node owning the destination address, herein called the destination node. Thus, at the point where a new connection is requested, its destination node will allocate to that connection a VLAN tag to use in combination with a destination address corresponding to that node. This VLAN tag will be allocated such that no existing connection to the destination address whose route crosses with & subsequently diverges from the new connection shall share a VLAN tag. This is such that where differential forwarding is required (at the divergence point), the destination address/VLAN tag pair of the diverging connections are distinct.

10

Additional constraints can be placed on the allocation as described elsewhere to improve pathological behaviour in case of inadvertent partial route removal in a network where broadcast-on-unknown is not fully disabled on all VLAN tags used for connections.

20

Alternately, for simplicity, but at a cost of reduced scalability in VLAN tag usage, the destination node shall allocate a unique VLAN tag for each connection going to a given destination address.

25

The VLAN tag having been allocated, it should be applied to packets where connections first do, or may, diverge in route and/or where the traffic first enters a traffic-engineered domain. Where encapsulation is not being performed, this would usually be where the traffic is first segregated by connection, for example by layer 3-7 filtering. Where the preferred method of encapsulation of customer data is being performed, the node containing the

- 22 -

encapsulation point is the preferred logical location for the application of the VLAN tag.

At a similar point towards the egress of the connection, the VLAN tag should be removed. Assuming the destination node is identified as the correct point to remove the VLAN tag allocate above, it is permitted for the destination node to use the destination address/VLAN tag pair, possibly in combination with the ingress port at which the packet arrives, as context for the interpretation of the payload.

It is helpful to identify the logical notion of the customer attachment interface. A customer attachment interface terminates each endpoint of a particular service to a particular customer. There may be one or many customer attachment interfaces associated with a port, or several ports. These interfaces may correspond to different services relating to the same customer or the same / different services relating to different customers. For ingress traffic at those ports, the function of the attachment interface is most sophisticated, as it is required to identify which incoming packets correspond to its associated service.

Those skilled in the art will realise that there are innumerable ways in which the packets to be transported through a given connection can be identified, and that this decision is a purely local function that occurs at the source node (with reference to the decapsulation/untagging destination) and so can be different among the plurality of tag-applying nodes. However, a few examples of the customer attachment interface/service/connection association may help to clarify:

- An association between a specific physical customer-facing port, an encapsulation function, an encapsulation service ID and a given connection.
- An association between a specific physical customer-facing port, a customer VLAN tag, an encapsulation function, an encapsulation service ID and a given connection, and a policer set to identify an acceptable quantity of traffic.

- 23 -

- A layer 3 - layer 7 switch identifying a SIP-signalled VoIP flow and establishing a new connection for it. The switch will then establish a suitable filter to identify all of the packets of the flow and attach the appropriate VLAN tag to its packets to route them along the established connection.
- 5 - An in-box logical association between a certain connection, encapsulation function and encapsulation service ID with a port of a virtual customer-address-space Ethernet switch instance (as per later VPLS explanation).
- An association between a physical port of a node (the service implicitly including all packets received from that port, with known destination MAC
- 10 addresses), a VLAN tag application function at that node, and a given connection.

The term 'service instance identifier' is used herein to define the identifier which identifies the customer attachment interface, which might also be referred to as an encapsulation service ID or alternatively as extended service
15 VLAN ID or service instance identifier (as in the current draft of IEEE 802.1ah).

Figure 6 demonstrates the actions on receiving an Ethernet frame (step 60) at an Ethernet switch after the stage at which VLAN tag has been assigned as described above, where the switch first selects a forwarding table based on the VLAN tag contained in the frame (step 62). Then, the switch selects an output port based on the destination address contained in the frame (step 64). Finally, the switch forwards the frame on the selected output port (step 66).

This method of differential forwarding using the combination of VLAN tag and destination address should be contrasted with the method of the prior art.
25 IEEE 802.1Q is exemplar of the prior art in which a VLAN is defined to be a partitioning of network resources. For example where those network resources may be defined in terms of ports. A VLAN is a specific set of ports, a subset of all the ports in the network. More specifically that subset of ports would be connected such that data frames may flow only between any of
30 those ports within the subset, and no others of the network. As a direct

- 24 -

consequence of this arrangement, any other subset of ports of the network disjoint from the first i.e. not connected must have a distinct VLAN tag. Whereas within the method of the present invention all ports of the network even where in disjoint subsets may have the same VLAN tag. This is because
5 the partitioning of network resources is achieved by the combination of VLAN tag and destination MAC address. The procedure by which the network is configured in the present invention to enable differential forwarding of connections has been described above.

The prior art of VLAN-aware bridges described above with reference to Figure
10 1 implement a structural mechanism by which VLAN tag is used first to determine the validity of the tag, and then to access a forwarding table to determine from the MAC address how to forward the frame. This same structural implementation may be reused by the present invention to functionally behave such that the combination of VLAN tag and destination
15 MAC address determine a differential forwarding behaviour. The advantage of using the same structural implementation, although not absolutely necessary, allows existing hardware Ethernet switches to be reused. However the means and rules by which the forwarding tables are populated is distinct from the prior art: according to the present invention, VLAN tags and entries in
20 forwarding tables corresponding to connections to be established across the carrier network are directly configured into the appropriate Ethernet switches using the connection setup process described above. In the preferred embodiment in which encapsulation is used, data traffic is associated with a particular connection on entry into the carrier network (more specifically at the ingress PE-Core) by giving the frames a selected VLAN tag as well as destination address (ie the MAC address of the egress PE-Core). Encapsulation in this context will ensure that the raw Ethernet frames
25 received from the customer will not be altered in this process.

Figures 7 and 8 show how the use of a combination of VLAN tag and
30 destination address may be used to differentiate connections. Figure 9 shows how a deliberate lack of differentiation in the combination of VLAN tag

- 25 -

and destination address and selection of port causes convergence of otherwise individual connections. Each of Figures 7 to 9 show connections across a carrier network comprising 4 provider edge Ethernet switches 71, 72, 73 and 74 (corresponding to PE1, PE2, PE3, PE4), further Ethernet switches in core 78 including core Ethernet switch 75, and communications links between the core and edge switches (reference numerals omitted for clarity).

In Figure 7, connections 76 and 77 have both the same source address (edge Ethernet switch 71 - PE1) and destination address (edge Ethernet switch 73 – PE3). However, the routes that connections 76 and 77 traverse are different. In particular, it can be seen that at core Ethernet switch 75, connections 76 and 77 converge and then immediately diverge. Despite the common destination address, core Ethernet switch 75 is able to differentiate frames belonging to connection 76 from frames belonging to connection 77 (and to forward them accordingly) on the basis of their different VLAN tags. Thus, data traffic in connection 76 has the VLAN tag 2, for example, whereas data traffic in connection 77 has the VLAN tag 1.

In Figure 8, connections 80 and 82 have both the same source address (edge Ethernet switch 71 - PE1) and are given the same VLAN tag (in this case the VLAN tag is 1), but have different destination addresses (connection 80 has edge Ethernet switch 73 – PE3 while connection 82 has edge Ethernet switch 74 – PE4). Again, the routes that connections 80 and 82 traverse are different. In particular, it can be seen that at core Ethernet switch 75, connections 80 and 82 converge and then follow the same path before diverging towards their destination points. Despite the common VLAN tags, core Ethernet switch 75 is able to differentiate frames belonging to connection 76 from frames belonging to connection 77 (and to forward them accordingly) on the basis of their different destination addresses.

From Figures 7 and 8 it can be seen that, differentiation between Ethernet frames belonging to different connections is achieved according to the

- 26 -

combination of destination address and VLAN tag. A difference in either may be used to achieve differential forwarding required for connections.

Figure 9 shows how a deliberate lack of differentiation in the combination of VLAN tag and destination address and selection of port causes convergence of otherwise individual connections. In Figure 9, connections 90 and 92 have the same destination address (edge Ethernet switch 73 - PE3), and are given the same VLAN tag (in this case the VLAN tag is 1), but have different source address (connection 90 has edge Ethernet switch 71 - PE1 while connection 92 has edge Ethernet switch 72 – PE2). Again, the routes that connections 90 and 92 traverse are different, but this is only because the data traffic is injected into the carrier network from different ingress points – ie edge Ethernet switches 71 and 72. Once the routes converge at core Ethernet switch 75, they stay converged until their destination at edge Ethernet switch 73. This is because they have the same destination address and VLAN tag and there is no way of differentiating them on the basis of the combination of destination address and VLAN tag alone.

The ability to establish a plurality of connections between any two points in the network has advantages of resiliency but also for in service maintenance. The is an aspect of the present invention of being able to arrange “make before break” connections in order to change transport characteristics, route around planned maintenance etc. The path re-route may be global with respect to the network, i.e. from source edge to destination edge, or may be local in the sense that a part of the path between any two given nodes on the original path are re-routed.

The actual means by which traffic may be switched in this manner is advantageously only required to change a forwarding rule at a single node for any given direction of a connection. A data traffic flow may be re-routed by simply provisioning a new connection with a different VLAN tag and then using that VLAN tag in the MAC header of the Ethernet frames at the ingress point of the original connection. Re-routing of data flows in this way does not

- 27 -

lose any frames since the new connection may be established contemporaneously with the old connection and new Ethernet frames directed into the new connection while earlier frames are still in transit over the old connection. Alternatively, the same VLAN tag and destination MAC address 5 may be configured in nodes not on the existing path but arranged such that the first and last such nodes are connected directly to nodes on the existing path. Then by changing only the port on the node of the existing path that connects to the first such new node, all subsequently transmitted frames are forwarded over the new path. The frames of the new and old paths may be 10 arranged to merge as per Figure 9 further downstream. Alternatively again, connections may be pre-established – ie in advance of any detected failure – over protection paths. Thus, re-routing data flows in the event of failure is even simpler and quicker, since the ingress point may immediately start transmitting frames over the pre-established connection. Advantageously, 15 and unlike conventional circuit-switched networks, pre-established connections take up no bandwidth until actually being used. Thus 1:n or 1:1 protection schemes may be implemented without resulting in inefficient network resource utilization. If working paths have a differential path delay compared to the protection paths, then buffering may be performed at one or 20 more convergence points between the protection and working path to avoid out-of-order frame delivery due to re-routing between working and protection paths. Sequence identifiers may be used in addition or alternatively to enable frame re-ordering, for example In this manner it is possible to effect a fully 25 error free (hitless) switchover.

A further advantage of connections in an Ethernet network is the ability to grow the network organically, in that new nodes and links may be introduced to the network without causing any interference with existing nodes, links, 30 connections or traffic. Any new connections may be established through the additional resources. Furthermore existing traffic may be switched on to the new resources by any of the methods described heretofore. Consequently

- 28 -

there is enormous flexibility in the management of real and pragmatic networks.

So far, only the establishment of point-to-point (ie unicast) connections and multipoint-to-point in the context of Figure 9 where the traffic is either merged or multiplexed (the traffic only meets at the egress node) have been described. However, according to the present invention, point-to-multipoint or multipoint-to-multipoint connections may also be established across Ethernet networks as will now be described. Conventional Ethernet switches are capable of a multicast service. Typically this is achieved by configuring the forwarding table with more than one output port (but not necessarily all output ports) for a given multicast destination address. According to the present invention, for relatively small scale multicast operation, a point-to-multipoint connection may be configured as described above but using a combination of VLAN tag and multicast address mapping to more than one output port (but not necessarily all output ports) of selected Ethernet switches. However, this approach is only suitable for relatively small scale multicast operation.

According to the present invention, a carrier network supporting a large number of point-to-multipoint or multipoint-to-multipoint connections, could be configured as a Resilient Packet Ring (RPR) emulated over the Ethernet MAC addressed network using multiple unicast connections established as described above. The following description is given in the context of a virtual private network (VPN) service, i.e. where there is a limited community of interest for each data frame. Two modes of operation are envisaged: a sparse mode for many customers with few sites, and a dense mode for few customers with many sites. The detailed mechanisms are described in one of the Applicants' co-pending US Patent Application Serial Number 10/698,833 (Nortel Networks Reference 15877RO) entitled Virtual Private Networks Within A Packet Network Having A Mesh Topology which document is incorporated herein by reference. The dense and sparse modes of operation will now be briefly described with reference to Figures 10 and 11.

- 29 -

Figure 10 shows a sparse mode of broadcast operation for many customers with few sites. Figure 10 shows a part of carrier network 20 comprising a part of fully-meshed core network 100, PE-Core edge Ethernet switches 104 a to d and PE-Edge edge Ethernet switches 102. Broadcast traffic 106a is received at PE-Core switch 104b from a customer site. Note that this traffic is broadcast within the context of a particular customer VPN, but is multicast within the context of the carrier network as a whole. The traffic is encapsulated and placed onto an RPR emulated by 4 uni-directional connections 108 a to d. The four connections are established as point-to-point connections as described above. The traffic is forwarded across each connection in turn until it reaches the start point again at PE-Core switch 104b. On receipt of an encapsulated frame, each endpoint of the four connections determines whether to process the frame for distribution to the customer via PE-Edge edge Ethernet switches 102 to which it is connected. This is done on the basis of broadcast destination addresses contained in the frames, and the VPN membership of customer sites attached to these Ethernet switches. Processing the frames involves decapsulating them and replicating them as required to one or more of PE-Edge edge Ethernet switches 102. It can be seen that no bandwidth need be dedicated to broadcast traffic in the sparse mode of operation since the four point-to-point connections may be trunked – ie they may be used to carry non-broadcast data and other customer's data, whether broadcast or not.

Figure 11 shows a dense mode of broadcast operation for few customers with many sites. Figure 11 shows a part of carrier network 20 comprising a part of fully-meshed core network 100, PE-Core edge Ethernet switches 104 a to d and PE-Edge edge Ethernet switches 102 as with Figure 10. Broadcast traffic 110a is received at PE-Core switch 104b from a customer site. Note, as above, that this traffic is broadcast within the context of a particular customer VPN, but is multicast within the context of the carrier network as a whole. The traffic is encapsulated and forwarded over a uni-directional connection 110b to a core switch 116a. Uni-directional connection 110b may

- 30 -

be trunked. At core switch 116a, the traffic is forwarded in over a bi-directional RPR 112 emulated by connections between core switches 116 a to d using a bidirectional connection between each pair of adjacent nodes. The RPR is dedicated to a particular customer's broadcast traffic and is not trunked. This is achieved by using a unique VLAN tag for forwarding in the RPR.

The traffic is forwarded around RPR 112 to each of the core switches 116 a to d in one direction or the other, whichever is shortest for each respective core switch. Each core switch broadcasts the received frames over uni-directional connections 114 a so that each of PE-Core switches 104 a to d receives the traffic. Then, as with the sparse mode of broadcast operation described above, each PE-Core switch determines whether to process the frame for distribution to the customer via PE-Edge edge Ethernet switches 102 to which it is connected. This is done on the basis of broadcast destination addresses contained in the frames and involves decapsulating and replicating them as required to one or more of PE-Edge switches 102 for onward transmission to the customer sites.

Figures 12 to 14 show exemplary arrangements of how Virtual Private LAN Services (VPLSs) may be provided according to the present invention. In VPLSs, potentially geographically distant hosts are provided with any to any communications capability over a carrier network which appears to the hosts to function as if the hosts were directly connected to a private Local Area Network (LAN). According to the present invention, VPLSs are provided over a meshed Ethernet carrier network over which connections may be established as described above. In each Figure, carrier network cloud 20 is comprised of Ethernet switches (only Ethernet switches located at the network edge are shown for clarity).

In Figure 12, one exemplary arrangement is shown in which 4 geographically distant customer sites (not shown) are respectively connected to 4 carrier edge nodes 120, 122, 124, and 126 which are themselves connected in a full

- 31 -

mesh via connections 130, established over the carrier network in accordance with the present invention, to provide a VPLS. In this arrangement, each of the carrier edge nodes 120, 122, 124 and 126 provides conventional Ethernet functionality associated with the physical port used for customer attachment to the PE-core. However, the Broadcast-on-Unknown function, MAC learning of customer addresses reachable via remote PE-cores, etc, is not associated with a physical Ethernet port on the core network side, but with a mesh of point to point connections between participating PE-cores, set up in accordance with the present invention. Thus, broadcast is performed by packet replication at the ingress PE-core, then encapsulation and unicast in the carrier domain. Each customer MAC address as it is learned is associated with the carrier network address and VLAN tag which is used to reach the remote PE-core.

In Figure 13, another exemplary arrangement is shown in which 4 geographically distant customer sites (not shown) are respectively connected to 4 carrier edge nodes 120, 122, 124, and 126 which are themselves connected in a hub and spoke arrangement via connections 132, 134 and 136, established over the carrier network in accordance with the present invention, to provide a VPLS. A switch/router connected to carrier edge node 120 acts as the hub whereas hosts or switches/routers respectively connected to carrier edge nodes 122, 124, and 126 act as spokes. This switch/router could be owned by the customer, or could be owned by the carrier, located in a central office, and used to offer service to more than one customer. This switch/router is the single entity in the carrier network which needs awareness of customer MAC addresses. Exactly as above, each customer MAC address as it is learned is associated with the carrier network address and VLAN tag which is used to reach the remote PE-core. If a single physical switch is used to support multiple customers, conventional VLAN technology can be used to ensure customer separation in the switch/router. The hub switch/router is responsible not only for providing communications between hosts connected to it and hosts connected to other customer sites,

- 32 -

but also for forwarding communications between hosts connected to any two other customer sites. In the latter case, traffic traverses two connections over the carrier network – for example, connection 132 from carrier edge node 122 to carrier edge node 120 and connection 134 from carrier edge node 120 to carrier edge node 124. At carrier edge node 120, after traversing connection 132, the traffic emerges from the carrier network. The hub switch/router identifies that the traffic is for another customer site and switches/routes it back onto the carrier network via the connection 134.

In Figure 14, another exemplary arrangement is shown in which 6 geographically distant customer sites (not shown) are respectively connected to 6 carrier edge nodes 140, 142, 144, 146, 148 and 150 which are themselves connected in a dual hub and spoke arrangement via connections 138, established over the carrier network in accordance with the present invention, to provide a VPLS. Switches/routers connected to carrier edge nodes 140 and 142 both act as hubs whereas hosts or switches/routers respectively connected to carrier edge nodes 144, 146, 148 and 150 act as spokes. As with the arrangement in Figure 13, the hub switches/routers are responsible not only for providing communications between hosts connected to them and hosts connected to other customer sites, but also for forwarding communications between hosts connected to any two other customer sites. In the latter case, as with the arrangement in Figure 13, traffic may traverse two connections or three connections if the customer sites are not directly connected by single connections to a single hub switch/router.

It will be appreciated that other arrangements of VPLSs are possible having any number of customer sites connected in a full mesh or in a single or multi hub and spoke arrangement or combinations of the two. The choice of arrangement will largely depend on the communities of interest in the various customer sites and the bandwidth required there between. In further embodiments, the two or multiple hub switches/routers may be provided at each customer site, each connected via connections to one or more other customer sites to provide load balancing and resiliency. The connections

- 33 -

used to provide VPLS may be unique to the customer or trunked in the sense that multiple customers use trunk connections.

Engineered connections set up according to the present invention may also be used to support engineered layer 3 any-to-any VPNs in accordance with
5 RFC 2547 when these are implemented as disclosed in [ref to 16898ROUS01U]. That disclosure describes a method and system for allowing an engineered virtual private networking solution through the use of a tandem routing device as a virtual hub in a logical hub and spoke network topology. This provides an engineerable alternative to the full mesh
10 connectivity between PE-cores normally used to support any-to-any services.

It will be apparent to those skilled in the art that that the designation for each VPN of a tandem switching point at layer 3 constrains the traffic flows from all PE-cores supporting a specific VPN to pass through said tandem switch, and so setting up connections according to the present invention between all PE-cores and the tandem switch allows the determinism and predictability at
15 layer 3, as described in co-pending US Patent Application 10/910,685 filed 4 August 2007, to be exploited in the transport layer as well.

Data plane monitoring functions such as heartbeat, ping and connection monitoring using the ITU-T standard-in-preparation currently referred to as
20 Y.17ethOAM has been mentioned above. These methods may be used to perform end to end connection monitoring and fault detection as follows. In one embodiment, a defined and well-known EtherType is assigned to Operations and Maintenance (OAM) frames, such as heartbeat or ping messages, which are sent across the end to end connections established in
25 the network using the same destination address and VLAN tag as data plane traffic is sent across those connections. The EtherType field is not used for forwarding in the network but is used to filter OAM traffic from data frames at the network edge. OAM frames may then be forwarded to OAM systems. Thus, OAM frames, such as heartbeat or ping messages, will be forwarded in
30 the network in the same way as of data plane frames and will therefore share

- 34 -

the same fate as data plane frames. Thus, OAM traffic can be used for fault detection and end to end connection monitoring. In a second embodiment, OAM traffic is not forwarded over the same connections as data plane traffic. Rather, separate OAM connections are established using one or more different destination addresses to those connections to be used for data plane traffic. Individual OAM connections are at least partially co-routed and possibly fully co-routed with individual data plane connections so that OAM frames share the same fate as data frames at least over the co-routed portions. However, the different destination addresses correspond to the address or addresses of a OAM management system or server. Advantageously, this avoids the need for line-rate filtering on EtherType at network edge nodes. It is also advantageous in a further embodiment to preferentially arrange a largely separate mesh of OAM connections across the network which can be monitored and used to "triangulate" the location of faults through correlation. This could be used to determine the affected data plane connections, while there may be little fate sharing between any one individual OAM connection and any one individual data plane connections.

The frequency of the ping or heartbeat frames may be used to adjust the distribution of the time interval to detect faults. The detection of faults may be used to trigger a protection switch on to a suitably configured protection path, and the trigger may be effected directly in hardware or software. Alternatively, the alarms may be filtered in a higher order system before triggering a protection switch to improve control of the network. End to end OAM in a connection may limit the frequency with which pings and heartbeat may be sent so as not to overload the network with this traffic type. In a further embodiment it is possible to implement pings and heartbeats on each individual link between Ethernet switches, or on segments of the end to end connections. Any detected failure on such a link or segment is straightforwardly correlated to the connections it affects by direct consultation of the forwarding tables in the Ethernet switches at either end. The alarm state may be propagated along all or some of the respective connections to

- 35 -

the end point to trigger the same protection mechanism above. However, the link or segment heartbeats may be sent at higher frequency than those of an end to end connection with a much lower chance by orders of magnitude of overloading the network, since links and segments are much fewer in
5 number.

To enforce QoS requirements, such as bandwidth guarantees, over connections established according to the present invention, admission control and policy enforcement may be implemented at ingress nodes as described above. Admission control may also be performed in the control or
10 management planes also as described above. Different classes of service may be provided for traffic forwarding by classifying customer traffic on the basis of customer one or more of the following: VLAN tag, IEEE 802.1p priority level, DiffServ codepoint (DSCP), MPLS EXP bits and so on. The classification may be mapped to carrier VLAN tag or using IEEE 802.1p or
15 802.1ad strict priority level, for example, for marking and segregation in the carrier network. Those skilled in the art will appreciate that classes of service may be distinguished in very many ways that are compatible with the present invention, which may be based on a mapping of a explicit marking of an incoming frame by a wide variety of possible fields to an explicit marking in
20 network 20, an explicit marking of an incoming frame by a wide variety of possible fields to an implicit forwarding treatment for a given connection in network 20, an implicit classification of a variety of possible fields in an incoming frame to an explicit marking in network 20, and an implicit classification of a variety of possible fields in an incoming frame to an implicit
25 forwarding treatment for a given connection in network 20. Those skilled in the art will also appreciate that an original marking may be remapped or remarked on egress by those frames from network 20. Conflicts for forwarding resources at nodes of the carrier network may be resolved by 1) using a strict priority scheduling scheme (such as IEEE 802.1p) in which frames of higher priority are always forwarded in preference to frames of lower priority; 2) using a weighted fair queuing scheduling scheme in which
30

- 36 -

classes of frames having lower priority still have some proportion of forwarding resources albeit lower than classes of frames having higher priority; or 3) using a differential discard eligibility scheduling mechanism in which the discard rate (a function of queue fill) applied to frames entering the queue for output over an output port of an Ethernet switch is different for different classes of traffic. In 3) above, the discard rate for classes of traffic having a lower priority is higher than the discard rate for classes of traffic having a higher priority for equivalent queue fill. Thus, proportionately more lower priority frames are discarded than higher priority frames as the output queue fills up. However, unlike in 1) and 2) above, frame disordering between different classes of traffic cannot occur because there is effectively only one queue. This has the advantage of permitting customers with guaranteed bandwidth connections to burst over agreed bandwidth limits using lower priority classes of traffic, without introducing potential disordering of frames. Those skilled in the art will appreciate that any or all of the mechanisms of classifying, marking, remarking, policing and scheduling may be applied to traffic according to the ability to differentiate connection forwarding using the combination of VLAN tag and destination MAC address, and any other fields of the frame as necessary.

It has been described above how connections may be established over a meshed Ethernet carrier network through configuring forwarding tables in network nodes and how data may be forwarded over those connections. The reader will appreciate that connections may be removed by deleting the configuration data from every node over which the connection was established. It is important that all such configuration data is removed to avoid network failure or inefficiency. The default behaviour of Ethernet switches on receiving a frame addressed to an unknown destination (ie where there is no forwarding state configured for that destination address) is to broadcast the frame out on all output ports. In simply-connected networks this behaviour is appropriate. However, with a meshed topology, this behaviour can be catastrophic. Through partial removal of connections (in

- 37 -

particular where configuration data is left at ingress points of a connection but deleted at points further along the connections towards or including the egress point), it remains possible that Ethernet frames for the PE may enter the network but arrive at a point where there is no configuration data for 5 forwarding them, resulting in undesirable broadcast behaviour. Furthermore, partial removal of connections may leave forwarding loops configured by accident.

One solution to the problem of partial removal of connections is to alter the behaviour of the Ethernet switches forming the carrier network so that instead 10 of broadcasting unknown traffic, they discard packets and possibly issue an alarm, log or count the discarded packets. However, altering the basic behaviour of Ethernet switches may require a hardware modification. While possible, this is not preferable. However, conventional Ethernet switches generally provide a software configurable function called rate limitation. 15 Preferably, at all or most switches of the carrier network rate limitation is used to set a rate of zero, or a low rate if zero is not possible, for broadcast traffic including broadcast-on-unknown traffic.

Where this is not possible, other pre-emptive approaches to minimising the problems of partial removal of connections may be used. One approach is to 20 use block lists otherwise known as access control lists or ACLs. Conventional Ethernet switches provide a block list (typically of limited length) which may be used to specify certain destination MAC addresses such that received Ethernet frames addressed to these blocked address will be discarded without forwarding. By blocking, at all or most nodes of the network, the MAC 25 addresses of many (but not all) MAC addresses of provider edge nodes it is possible to minimise the potential dangers of partial removal of connections without over restricting the carrier's flexibility in establishing connections across the network. Notably, it is necessary to block different MAC address at different nodes of the network. Typically, at a given node, the block list will 30 include only the MAC address for provider edge nodes to which no

- 38 -

connections are likely to be established through that node. This approach is not easily scaleable with large networks (the limited number of entries in block lists may be exhausted by large numbers of provider edge nodes). However, note that to prevent loops it is only necessary to block rogue frames at one node in any loop. Thus, it is possible to "spread" the blocked destination addresses more thinly across the network and still provide a degree of protection from loops thereby making more efficient use of the limited capacity of block lists.

While it is the use of VLAN tags in the present invention that enables flexibility in establishing connections across the network, the failure to remove connection state fully leaves the potential for looping of traffic. In particular, the problem will arise where a logical loop is left configured for any single given VLAN tag – ie the output ports of nodes defining a physical loop are left configured with membership of any single VLAN tag. Thus, another pre-emptive approach to minimising the problems of partial removal of connections is to allocate connections to or from neighbouring or nearby provider edge nodes using mutually exclusive VLAN tag pools. Thus, for example all connections to or from provider edge node PE1 will be guaranteed to have a different VLAN tag to those to or from neighbouring provider edge node PE2. In this way, loops including both PE1 and PE2 cannot accidentally be formed through the partial removal of connections since by definition any state left configured in PE1 and PE2 will use different VLAN tags. This approach may be generalised by allocating connections to or from n adjacent provider edge nodes using n mutually exclusive VLAN tag pools. n is chosen to be sufficiently large to segregate use of VLAN tag pools as much as possible while providing sufficient flexibility in connection establishment to or from any particular provider edge node (bearing in mind that there are only 4094 possible VLAN tags). With smaller carrier networks it may be possible for each provider edge node to use a different VLAN tag pool. However, with larger carrier networks it will be necessary to re-use VLAN tag pools at topologically distant provider edge nodes otherwise

- 39 -

flexibility in connection establishment will be compromised though VLAN tag pools being too small. A further embodiment that does not require manual administration is for each Ethernet device to pick a random starting point in the VLAN range for allocation.

5 It will be appreciated that combinations of the above approaches to minimising the problems of partial removal of connections may be employed.

Mixed-mode networks so called because they combine conventional Ethernet, 802.1Q, or other forwarding modes, simultaneously with the connection-oriented forwarding mode present invention, may also be implemented. In mixed-mode networks, part of the VLAN tag space (for example VLAN tags 1-2048) is assigned to conventional mode Ethernet forwarding and operates using a VLAN-aware spanning tree protocol and auto address learning. Another part of the VLAN tag space (for example VLAN tags 2049-4096) is assigned to connection-oriented mode Ethernet forwarding as described above. Note that the use of two or more such VLAN spaces creates logically separate forwarding mode networks over the same physical network. Forwarding state installed in Ethernet switches in one mode (for example through auto address learning and spanning tree) is differentiable from forwarding state installed in Ethernet switches in another mode (for example through control plane connection setup procedures according to the present invention) by having different VLAN tag spaces assigned. Thus, forwarding state in one mode, and mechanisms for installing or removing such forwarding state, do not affect forwarding of traffic in another mode and vice versa.

25 In mixed-mode Ethernet networks, preferably the connection-oriented Ethernet mode is given a higher forwarding priority than the conventional Ethernet mode so that QoS (in particular bandwidth and latency) may be guaranteed for connection-oriented Ethernet mode connections. This may be achieved through assigning a higher priority level to frames belonging to the

- 40 -

connection-oriented Ethernet mode VLAN space using IEEE 802.1p strict priority, for example. Conventional Ethernet mode forwarding may then be offered as a best efforts service and may be used to allow customers guaranteed connections to burst over agreed bandwidth guarantees when overall network loading permits. Another possible example of a mixed-mode Ethernet network would be to have one part of the VLAN tag space (for example VLAN tags 1-1024) assigned to conventional mode Ethernet forwarding, another part (for example VLAN tags 1025-2048) assigned to VLAN label switching (as described in draft-kawakami-mpls-lsp-vlan-00.txt, for example) and another part (for example VLAN tags 2049-4096) assigned to connection-oriented mode Ethernet forwarding as described above. Advantageously, control plane metrics may be "leaked" (ie purposefully communicated) between different forwarding modes so that, for example, spanning tree converges on virtual topologies that will avoid heavily loaded links primarily utilizing the connection-oriented mode. This is in addition to use of priority above to mitigate resource collisions between modes.

While embodiments have been described above with reference to the use of VLAN tags for enabling flexibility in establishing and differential forwarding of data traffic associated with different connections, the reader will appreciate that other tags or identifiers may be used. For example, MPLS labels may be used. In this case, the MPLS labels are appended, prepended or inserted into the Ethernet frames and Ethernet switches in the network forward based on a combination of Ethernet destination address and MPLS label. Note, that this is entirely different to conventional use of MPLS labels since the MPLS labels are not used for label switching.

Also, while embodiments have been described above with reference to Ethernet networks and Ethernet frames, those skilled in the art will appreciate that the present invention applies in general to any frame-based, packet-based or cell-based switching network whether at OSI layer 2 or layer

- 41 -

3 network. And to data structures including frames, packets and cells. In the following claims, the term frame-based network, or cognate terms, shall denote any such switching network and the term frame, or cognate terms, shall denote any such data structure. For example, IP networks comprising a
5 mesh of IP routers may be used to route IP packets. Conventional Open Shortest Path Forwarding (OSPF) control plane mechanisms would be disabled to allow direct configuration of forwarding or routing tables. In this case, the routers may be configured to route on a combination of IP destination address and VLAN tags, MPLS labels, DiffServ codepoints, IPv6
10 flow labels, type of service, traffic class or other such fields, or optional fields added specifically to act as identifiers. This is of particular interest where IP/IP, IP/MPLS and Psuedo Wire/IP or similar forms of encapsulation are used in order to maintain customer separation over this forwarding layer.

15 It will also be appreciated that addresses other than destination address may be used in combination with a qualifying identifier to enable differential forwarding according to the present invention and thereby to enable the establishment of connections. In particular, forwarding may be performed on the basis of a source address of data frame corresponding to a source node.
20 Furthermore, forwarding may be performed on the basis of any address field contained in a data frame or specifically added to a data frame. Furthermore, differential forwarding may be performed on the basis of a combination of source and destination address, or on the basis of a single address field which can uniquely address a sufficiently large space of nodes and additionally can qualify the address to enable differential forwarding of data frames.
25 The reader will appreciate that the methods described above may be implemented in the form of hardware or software operating on conventional data processing hardware.